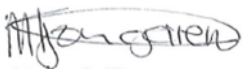
Bank of Namibia

No.: 6/2026 **2026**

GUIDELINES UNDER THE BANKING INSTITUTIONS ACT, 2023 (ACT NO. 13 OF 2023): USE OF ARTIFICIAL INTELLIGENCE IN THE BANKING SECTOR

In my capacity as Director of Banking Supervision and under the powers vested in the Bank under section 3(1) of the Banking Institutions Act, 2023 (Act No. 13 of 2023), I hereby issue the **Guidelines on the use of Artificial Intelligence in the Banking Sector (BIG-6/2026).**



MAGRETH TJONGARERO
DIRECTOR: BANKING SUPERVISION **28 OCTOBER 2026**
WINDHOEK

TABLE OF CONTENTS

PART I: PRELIMINARY 2
PART II: STATEMENT OF POLICY 5
PART III: IMPLEMENTATION AND SPECIFIC REQUIREMENTS 7
8. GOVERNANCE AND ACCOUNTABILITY 7
9. HUMAN OVERSIGHT 10
10. AI INVENTORY AND RISK TIERING 12
11. DATA GOVERNANCE AND QUALITY 14
12. TRAINING, TESTING AND MONITORING OF AI SYSTEMS 16
13. EXPLAINABILITY AND INTERPRETABILITY OF AI SYSTEMS 17
14. SECURITY, DATA PRIVACY, CUSTOMER CONSENT AND PROTECTION 18
15. ETHICAL USE OF AI SYSTEMS 18
16. SUPERVISORY NOTIFICATION AND INFORMATION REQUIREMENTS 19

Public

PART I: PRELIMINARY

1. **Short Title** – AI Guidelines.

2. **Authorisation** – These Guidelines are issued in terms of section 3(1) of the Banking Institutions Act, 2023 (Act No. 13 of 2023).

3. **Application** – These Guidelines apply on a comply-or-explain basis and set out the Bank's supervisory expectations regarding the governance and use of AI in the banking sector. Banking institutions are expected to implement the principles and expectations contained in these Guidelines in a manner proportionate to their nature, size, complexity, risk profile and level of AI adoption. Where a banking institution adopts alternative measures or is unable to fully implement a particular expectation, it should be able to demonstrate to the Bank that equivalent or more effective outcomes are achieved and that implementation is progressing in a reasonable and risk-appropriate manner.

   These Guidelines are not legally binding. References to supervisory expectations and controls should therefore be interpreted in accordance with the principles-based and proportionate nature of the Guidelines. The Bank will assess implementation based on the effectiveness of a banking institution's governance, risk management and control environment, and the outcomes achieved through the application of these Guidelines.

4. **Acronyms:**

4.1	AI	Artificial Intelligence
4.2	AIS Program	AI Systems Program
4.3	BID-30	Determination on Information Security
4.4	BID-34	Determination on Outsourcing
4.5	ISO/IEC	International Organisation for Standardisation / International Electrotechnical Commission
4.6	GenAI	Generative AI
4.7	NIST	The National Institute of Standards and Technology

5. **Definitions** – Terms used in these Guidelines will have the same meaning assigned to them in the Act, as further defined below, or as reasonably implied by contextual usage. Where applicable, definitions are aligned with internationally recognised standards on AI, including relevant ISO/IEC standards on AI terminology and concepts to support consistency and clarity.

   5.1 **Act -** means the Banking Institutions Act, 2023 (Act No. 13 of 2023).

   5.2 **AI Incident** - means an event arising from the design, development, deployment, operation, performance, outputs, modification, failure, malfunction, misuse, or unexpected behaviour of an AI system, including circumstances where the AI system operates outside its intended purpose, parameters, or performance thresholds, that results in, or has the potential to result in, adverse impacts on customers, employees, critical operations, regulatory compliance, or the safety and soundness of the

Public

institution. For the avoidance of doubt, an underlying business, fraud, cyber, operational or other incident does not constitute an AI Incident solely because an AI system was used to monitor, detect, analyse, recommend or respond to the underlying event. Such an incident should be regarded as an AI Incident where the AI system itself caused, materially contributed to or exacerbated the adverse outcome.

5.3 **AI System** - means a machine-based system that, for explicit or implicit objectives, infers from input data how to generate outputs such as predictions, content, recommendations, or decisions that can influence real or virtual environments, and that can operate with varying levels of autonomy.

5.4 **Agentic AI** - means AI systems designed to plan and execute sequences of actions toward an assigned objective, which may include using tools or triggering automated workflows under defined guardrails.

5.5 **AI Governance Framework –** means a structured set of principles, policies, roles, processes, and controls established by an organisation to direct and oversee the responsible development and use of AI systems, in alignment with organisational objectives.

5.6 **AI Inventory -** means the institution's central register of AI systems in use or under development, including each system's owner, purpose, risk tier, approvals, and review status.

5.7 **AI Lifecycle -** means the stages of an AI system from design and data sourcing through development, validation, deployment, monitoring and decommissioning.

5.8 **AI Systems Program -** means a coordinated set of operational processes, resources, and lifecycle activities through which an organisation designs, develops, validates, deploys, monitors, and maintains AI systems in accordance with its AI Governance Framework. An AI Systems Program is contained within the AI Governance Framework.

5.9 **Artificial Intelligence –** means the capability of a system to perform tasks that typically require human intelligence, including learning, reasoning, problem-solving, and generating outputs based on data.

5.10 **Bank -** means the Bank of Namibia as referred to in section 2 of the Bank of Namibia Act, 2020 (Act No. 1 of 2020).

5.11 **Banking institution -** means a public company authorised under the Act to conduct banking business.

5.12 **Bias -** refers to systematic differences in the outputs or outcomes of an AI system that may result in unfair or unintended treatment of individuals or groups, arising from data, model design, or the operation of the system.

5.13 **Controlled autonomous operation** - means a mode of AI system operation whereby the system performs tasks and makes decisions autonomously within predefined boundaries, controls, and risk limits established by human authorities, with mechanisms in place to enable monitoring, intervention, and override where necessary.

5.14 **Data drift -** refers to changes in the distribution or characteristics of input data over time, which may affect the performance or reliability of an AI system.

5.15 **Data provenance** - means the documented origin and sourcing history of data used to train, validate, test, or operate an AI system.

Public

5.16 **Explainability** - means the extent to which the internal logic, behaviour and outputs of an AI system can be understood and articulated in a manner that is meaningful to relevant stakeholders, including users, decision-makers and supervisory authorities.

5.17 **Fairness** - means the design, development and use of AI systems in a manner that seeks to avoid unlawful discrimination, inappropriate bias and unjustified adverse outcomes, taking into account the nature, purpose and context of the AI system.

5.18 **Frontier AI Model** - means a highly capable General-Purpose AI Model that demonstrates advanced reasoning, content-generation, decision-making or autonomous capabilities and may present heightened or systemic risks due to its scale, capabilities or level of autonomy.

5.19 **General-Purpose AI Model** - means an AI model capable of performing a broad range of distinct tasks and that may be integrated into, or form the basis of, different AI systems and applications.

5.20 **Generative AI** - means AI models that generate content such as text, images, audio, video, or code in response to inputs.

5.21 **Human-in-the-loop -** means a form of human oversight whereby a human is actively involved in the decision-making process of an AI system, including the review, validation, and approval of outputs prior to implementation, with the authority to intervene, override, or stop the system's operation where necessary.

5.22 **Human-on-the-loop -** means a form of human oversight whereby a human monitors the operation and outputs of an AI system on an ongoing basis, without being directly involved in each decision, and retains the authority and capability to intervene, override, or halt the system's operation where necessary.

5.23 **Human overseer** – means a natural person responsible for exercising human oversight over an AI system, including monitoring, reviewing, challenging, escalating, overriding or stopping system outputs where necessary.

5.24 **Human oversight** - means the ability for human intervention in an AI system, including monitoring, reviewing, and, where necessary, challenging or overriding system outputs.

5.25 **Interpretability** - means the degree to which a human can directly understand the cause-and-effect relationships within an AI model, including how input variables influence outputs, without requiring additional explanation techniques.

5.26 **Machine Learning** – means a subset of artificial intelligence that uses computational techniques to enable a system to learn from data or experience and improve its performance on a specific task, or to make predictions or decisions, without being explicitly programmed for each outcome.

5.27 **Material change -** means a change to an AI system, its data, integration, intended purpose, operating environment or level of autonomy that may significantly affect the system's risk profile, performance, outputs, governance requirements or assigned risk tier. Such changes may include significant changes to training data, model architecture, decision-making logic, customer-facing functionality, system integrations, operating environments, levels of autonomy, or the deployment of an AI system for a new or materially different use case.

5.28 **Material incident -** means any event involving an AI system that results in or has the potential to result in a material adverse impact on customers, critical operations, safety,

Public

regulatory compliance, or the financial soundness of the banking institution, and that requires notification to the Bank under section 16 of these Guidelines. In determining whether an incident is material, banking institutions must consider all relevant circumstances, including the scale of impact, number of customers affected, extent of operational disruption, regulatory implications, whether sensitive data has been compromised, and the criticality or risk tier of the affected AI system. The criticality or risk tier of an AI system should inform, but not solely determine, the assessment of incident materiality.

5.29 **Model -** means the computational or statistical representation used by an AI system to generate outputs based on input data.

5.30 **Model drift -** refers to a decline in the performance or reliability of an AI model over time due to changes in input data, underlying relationships, or operating conditions compared to those present at the time of model training.

5.31 **Risk tier -** means the classification assigned to an AI system under the institution's risk-based methodology, for example, high, medium, or low, which determines minimum control and governance expectations in these Guidelines.

5.32 **Testing Data -** means data used to independently evaluate the performance, reliability, robustness and suitability of an AI system against defined requirements and acceptance criteria prior to deployment or following a material change.

5.33 **Training Data -** means data used for training an AI model and helping it develop algorithms.

5.34 **Transparency -** means the provision of information regarding the purpose, use, operation, limitations and governance of an AI system in a manner that is clear, understandable and appropriate to the needs of relevant stakeholders.

5.35 **Validation Data -** means data that provides an unbiased evaluation of a model fit on the training data set while tuning the model's hyperparameters.

PART II: STATEMENT OF POLICY

6. PURPOSE

These Guidelines set the minimum requirements for the safe, fair, ethical, explainable, transparent and accountable use of AI by banking institutions to protect customers, promote financial stability and maintain a stable and sound financial system. The Guidelines further support responsible innovation and controlled experimentation with AI.

These Guidelines establish a governance and risk management baseline that covers Board and Senior Management accountability, lifecycle controls from design through decommissioning, third-party oversight, data protection, and transparency to customers and supervisors.

These Guidelines are aligned with international standards and frameworks on AI, including ISO/IEC 42001 Artificial Intelligence Management Systems, ISO/IEC 23894 Artificial Intelligence Risk Management, and the NIST Artificial Intelligence Risk Management Framework. Banking institutions are expected to implement additional controls where necessary to ensure adequate

Public

risk mitigation beyond the minimum requirements set out in these Guidelines. The Guidelines complement existing regulatory and risk management frameworks and established international standards and should be applied alongside enterprise risk management, information security, and model risk policies rather than as a replacement.

The Bank recognises that banking institutions are at different stages of AI adoption, governance maturity, technological capability and operational complexity. Accordingly, the implementation of these Guidelines should reflect the nature, size, complexity, risk profile and level of AI maturity of the institution.

Given the evolving nature of AI technologies, industry practices and supervisory expectations, the Bank may review and amend these Guidelines from time to time. The Bank will apply a phased and proportionate supervisory approach and may determine, based on industry adoption, emerging risks and implementation outcomes, whether future supervisory measures, including Determinations or Regulations, are necessary.

The Guidelines are intended to support responsible innovation and experimentation with AI and should not be interpreted as prohibiting the adoption of emerging technologies, provided that risks are appropriately managed. Banking institutions should establish appropriate governance arrangements for AI experimentation, proof-of-concept activities, pilot deployments and controlled testing environments. Such arrangements should be commensurate with the nature, scale and risks of the activity and should support safe, responsible and controlled experimentation while maintaining sufficient flexibility to foster innovation and learning.

These Guidelines establish an initial supervisory baseline for the governance and oversight of AI in the banking sector, with the long-term objective of supporting the safe, fair, ethical, explainable, transparent and accountable adoption of AI while promoting innovation, financial stability and customer protection.

7. SCOPE

These Guidelines apply to all AI systems that generate outputs such as predictions, recommendations, content, or decisions that can influence outcomes in a banking institution's customer interactions, internal operations, or risk management processes. Covered systems include traditional machine learning models, general-purpose AI and GenAI, agentic AI, frontier AI models and future AI technologies of a similar nature or capability. The scope applies whether systems are developed in-house or sourced from third parties, and whether they are deployed on-premises or in cloud environments. Conventional rule-based systems, deterministic workflows, and standard statistical calculations operating solely on predefined instructions without AI capabilities fall outside the scope of these Guidelines.

The Guidelines apply proportionately to AI capabilities embedded within third-party applications, cloud services, operating platforms and outsourced arrangements. In determining the extent to which the requirements of these Guidelines apply, banking institutions should consider factors such as their level of reliance on the AI capability, the materiality of its impact on customers,

Public

operations or risk management processes, the assigned risk tier of the AI system, and the institution's level of visibility, control and ability to influence the operation, outputs or governance of the AI capability. Banking institutions remain responsible for ensuring an appropriate level of oversight, understanding, risk management and explainability commensurate with the risk tier, materiality and impact of the AI system.

Requirements apply across the banking institution, and where applicable, shared service centres and third-party arrangements to the extent that they support, affect or are relied upon in the banking institution’s operations in Namibia. They cover all functions that design, procure, integrate, deploy, operate, or oversee AI systems, including product, business operations, risk, compliance, data, technology, cybersecurity, and model risk management. Coverage applies to directors, employees, third-party service providers, and contractors involved in these activities.

The principles set out in these Guidelines must be implemented in a manner that reflects the nature, size, complexity and risk profile of the banking institution and the assigned risk tier of each AI system. Where terms such as appropriate, adequate, effective, timely, or regular are used, implementation must be assessed in the same context.

The Guidelines cover the full AI lifecycle. This includes design, data sourcing and preparation, training, validation, deployment, monitoring, change control, and decommissioning, with appropriate documentation and supporting evidence maintained at each stage. Such evidence may, where proportionate to the system’s risk tier, include documentation such as an AI impact assessment, pre-deployment validation or approval records, data sheet, validation report, monitoring plan, change log, incident log, customer disclosure records and decommissioning records.

PART III: IMPLEMENTATION AND SPECIFIC REQUIREMENTS

8. GOVERNANCE AND ACCOUNTABILITY

8.1 Role of the Board of Directors

a) The Board of Directors must:

i. take overall accountability for the oversight of AI governance, risks and outcomes, including where third-party or cloud-based AI services and foundation models are used.

ii. approve the institution-wide AI Governance Framework, including lifecycle controls from design and data sourcing through validation, deployment, monitoring, change control, and decommissioning.

Public

iii. approve a documented AI strategy, risk appetite and resource requirements aligned with institutional objectives, either as standalone documents or as part of broader Board-approved governance frameworks, which define permissible uses of AI, required levels of human oversight and explainability, and proportionality based on materiality and risk.

iv. review the AI strategy at least annually, and following any material changes in AI use cases, risk profile, operating environment or applicable regulatory requirements, to ensure that it remains appropriate and effective.

v. oversee the effective integration of AI risks into the institution's overall governance, risk management and control framework, and ensure that AI risks are managed through a three-lines-of-defence model, or an equivalent structure appropriate to the size, nature and complexity of the banking institution.

vi. ensure that the Board of Directors and Senior Management collectively possess sufficient understanding of AI-related opportunities and threats, and that the institution invests in appropriate training and skills to support effective oversight.

vii. ensure that responsibilities for AI risk management are clearly assigned to Senior Management, relevant Board Committees, and internal control functions for the management of AI risks, including the designation of accountable owners for each AI system.

viii. approve ethical principles and minimum standards to ensure non-discrimination, fairness, risk-appropriate explainability, and proportionate transparency and recourse for affected customers.

ix. ensure that oversight of AI systems, including validation, risk management, and internal audit functions, is independent from model development, deployment, and operation and that these functions are adequately resourced to challenge AI risks effectively.

8.2 Role of Senior Management

a) The reference to Senior Management in these Guidelines includes Principal and Executive Officers as approved by the Bank in terms of section 53 of the Act.

   Senior Management is responsible for implementing the AI Governance Framework, policies, standards, procedures and controls approved by the Board of Directors and for ensuring that AI systems operate within the institution's approved risk appetite and governance arrangements.

b) Senior Management must:

Public

i. implement the AI Systems Program and ensure that all business units apply lifecycle controls from design through decommissioning. Where AI systems are procured from third parties, the AI Systems Program should include appropriate due diligence, oversight, monitoring and risk management arrangements commensurate with the institution's reliance on, and level of control over, the AI system.

ii. establish and maintain a single institution-wide inventory or register of AI systems capturing ownership, purpose, data lineage and provenance, risk tier, validation and monitoring status, incidents and version history and ensure that proportionate controls are applied based on risk.

iii. establish and resource a cross-functional AI Governance function, which may be implemented through a dedicated AI Governance Committee or existing IT, technology, data, innovation or equivalent committees. The role of the function is to ensure effective review and approval of material AI use cases, track conditions and actions, and escalate matters to the Board or relevant Board Committee where appropriate.

iv. ensure that internal structures responsible for AI risk management are clearly defined and adequately resourced with appropriate skills, tools and capability.

v. promote a culture of AI risk awareness across the institution.

vi. integrate bias assessment and mitigation, risk-appropriate explainability and customer transparency and recourse into pre-deployment validation and ongoing monitoring processes.

vii. implement monitoring processes to track performance, fairness, and model drift on an ongoing basis. Such monitoring may be supported by appropriate tools, metrics, assessments, or other mechanisms commensurate with the nature, complexity, materiality and risk tier of the AI system.

viii. conduct risk-based due diligence on third-party providers and cloud providers whose products, services or platforms provide AI capabilities used by the banking institution, including assessment of their operational capability, security controls, data rights and location, training data quality and provenance, explainability, and bias safeguards.

ix. where AI systems are outsourced or provided by third-party service providers, banking institutions must ensure that outsourcing risks are appropriately identified, assessed and managed in accordance with the Determination on Outsourcing (BID-34).

Public

x. ensure that appropriate contingency planning and operational continuity arrangements are in place for outsourced AI systems, including the assessment of the service provider's ability to maintain business continuity, disaster recovery capabilities, and resilience of critical services, in terms of BID-34. The requirements of these Guidelines are intended to complement, and not replace, the outsourcing and insourcing requirements set out in BID-34.

xi. ensure that appropriate business continuity and disaster recovery arrangements are in place for AI systems, including fallback processes, failover mechanisms and periodic testing to ensure resilience in the event of system disruption or failure.

xii. maintain robust incident management processes and notify the Bank within 24 hours of becoming aware of a material incident, as defined in these Guidelines.

xiii. where materiality cannot be conclusively established immediately, banking institutions should exercise sound judgement and, where appropriate, notify the Bank based on the information available at the time.

xiv. report to the Board or Board Committee on AI-related risks and opportunities, and on the effectiveness and adequacy of the AI Governance Framework at a frequency commensurate with the level of risk.

xv. review the AI Governance Framework, policies, and procedures at least triennially and following any material changes in use cases, models, data or regulations to ensure that they remain relevant and effective.

xvi. develop and maintain role-specific training programs to support AI literacy across the institution, including for model owners, approvers, and oversight functions, and ensure that competency requirements are commensurate with the responsibilities of the relevant roles.

xvii. ensure that AI systems are adequately protected against data breaches, cyber threats, and unauthorised access, including AI-specific risks such as model manipulation and data poisoning, in terms of the requirements of BID-30.

9. HUMAN OVERSIGHT

9.1 Banking institutions must ensure that human oversight over AI systems is applied in a manner proportionate to the AI system's risk tier, throughout the AI lifecycle, to safeguard the integrity, fairness, explainability and lawfulness of decision-making.

9.2 Banking institutions must design AI systems to enable effective human oversight prior to deployment, including the definition of decision thresholds, alert mechanisms, escalation triggers, and stop conditions that are operationally practical and clearly understood by users.

Public

9.3 For each AI system, banking institutions must establish and document an oversight protocol that is proportionate to the system's risk profile. Oversight models may include:

a. Human-in-the-loop;
b. Human-on-the-loop; or
c. Controlled autonomous operation, subject to predefined guardrails and prior approval where required, including clearly defined limits on system actions, monitoring controls, and the ability for timely human intervention where necessary.

Banking institutions may determine the most appropriate oversight model, considering the AI system's risk tier, complexity, materiality and intended use.

9.4 High-risk AI systems must operate only under enhanced human oversight, including:

a. Clearly defined intervention rights;
b. Pre-approved operating limits; and
c. The ability to suspend or terminate the system's operation where necessary.

9.5 Where GenAI systems are used, particularly in customer-facing, legal, compliance, credit decision-making, or other high-impact processes, banking institutions must ensure that outputs are subject to appropriate human review and validation prior to use or reliance, commensurate with the system's risk and materiality.

9.6 Each high-risk AI system must have a clearly designated accountable owner or owners, supported by appropriate governance structures and roles, to ensure effective oversight and management of the system. Such arrangements must include clearly defined responsibilities across business, technical and operational oversight functions. The model owner or owners must be:

a. competent and appropriately trained; and
b. accountable for the system's operation, including the ability to interpret outputs, understand limitations, override decisions and stop the system where required.

9.7 Banking institutions must equip human overseers with appropriate tools, clear escalation paths, and defined response timelines to act promptly on alerts and exceptions.

9.8 During operation, banking institutions must monitor performance, bias, drift, and exceptions using appropriate monitoring approaches commensurate with the nature, complexity and use case of the AI system. Such approaches may include, where applicable, the use of dashboards, alerts or other monitoring tools, and should set clear

Public

triggers for human review, temporary pause or rollback where defined limits are breached.

9.9 Banking institutions must monitor AI systems for:

a. Performance degradation;
b. Bias and fairness risks;
c. Model drift; and
d. Operational exceptions.

9.10 Banking institutions must, in a manner proportionate to the AI system's risk tier, record, investigate and review incidents, overrides and interventions, including their outcomes, to support ongoing improvement of AI system controls and governance. For medium and high-risk AI systems, such records must be maintained in a structured and auditable manner.

9.11 Banking institutions must retain logs of alerts, overrides, stop events, and customer appeals together with version histories, in line with applicable records management requirements and internal policies. Retention periods should be commensurate with the system's risk tier and sufficient to support auditability, investigation and regulatory review.

10. AI INVENTORY AND RISK TIERING

10.1 Banking institutions must maintain a single, institution-wide view of all AI systems in use or under development within the banking institution, supported by a centrally governed AI inventory. AI inventory records may be maintained across different business units, functions or platforms, provided that a single authoritative institution-wide view is maintained. For purposes of the AI inventory, AI systems and use cases should be recorded once they have progressed to experimentation, proof-of-concept, development, testing or validation. Purely conceptual ideas and vendor demonstrations that have not progressed into a potential institutional use case need not be recorded in the AI inventory.

10.2 Each AI system must be recorded in the inventory with sufficient information to support effective governance, oversight and risk management. Such information may include system ownership, version, business purpose, risk tier, deployment status, human oversight approach, testing and validation status and key third-party dependencies, where applicable.

10.3 Where AI capabilities are embedded within broader third-party applications, platforms or tools, banking institutions may record the overall application, platform or tool in the AI inventory rather than each individual AI capability, taking into account the materiality of the AI capability and the institution's level of reliance, visibility and control.

Public

10.4 Banking institutions must adopt a documented risk tiering methodology that classifies AI systems as high, medium, or low risk using clear, objective criteria and defined thresholds.

10.5 Banking institutions may leverage existing risk classification, model risk management, technology risk management or similar governance frameworks for AI systems, provided they can demonstrate that such frameworks adequately identify, assess and manage AI-related risks and achieve outcomes consistent with the expectations of these Guidelines.

10.6 Where existing frameworks are utilised, institutions should ensure that AI-specific risks are appropriately considered and that governance, validation, monitoring and oversight measures remain commensurate with the level of risk presented by the AI system.

10.7 Risk tiering must consider factors such as the importance of the decision being supported, the potential impact on customers or systems, the level of system autonomy, whether outcomes can be reversed, the sensitivity of data being processed, the complexity of the model and the extent of reliance on third-party or external services.

10.8 At a minimum, high-risk classification should apply where AI systems influence access to essential banking services, have a material impact on customers or employees, or process sensitive personal information or involve a high degree of autonomy. Examples of high-risk AI systems may include credit decisioning models, AI-driven fraud detection systems with automated blocking actions, or customer profiling systems that materially affect outcomes.

10.9 Where a high-risk AI capability is embedded within a third-party application, platform or service, and the banking institution does not have direct access to the underlying model or technical information, the requirements of this section shall be implemented in a manner proportionate to the institution's level of access, visibility and control over the AI capability. Banking institutions may rely on available vendor documentation, independent assurance reports, contractual arrangements and output-based testing, where appropriate, to support compliance with these requirements.

10.10 High-risk AI systems should be subject to enhanced controls, including independent validation, heightened human oversight, rigorous monitoring, and governance approval, reflecting their higher risk level and potential impact.

10.11 Independent validation should be conducted by suitably qualified persons who are sufficiently independent from the design, development, deployment and operation of the AI system to provide objective assessment and effective challenge.

10.12 Medium-risk AI systems may include those supporting decision-making or customer engagement without directly determining outcomes, while low-risk AI systems may include routine or back-office automation with limited impact. Examples of medium-risk

Public

AI systems may include customer service chatbots, marketing and campaign management tools, productivity assistants and decision-support tools that do not directly determine customer outcomes.

10.13 Medium-risk systems should be subject to proportionate validation, oversight, and monitoring, while low-risk systems may be subject to simplified controls.

10.14 Banking institutions must determine and document the minimum control expectations for each AI system based on its assigned risk tier. Such control expectations should be proportionate to the assigned risk tier and materiality of the AI system, and may include testing depth, validation frequency, monitoring cadence, reporting frequency, and change approval levels.

10.15 High-risk AI systems must be subject to enhanced governance, validation, explainability, monitoring and reporting requirements, while medium- and low-risk AI systems may be subject to simplified or reduced controls, commensurate with their risk and materiality.

10.16 The AI inventory must clearly flag critical third-party AI services and record their substitutability, key dependencies, and exit options.

10.17 Where practical, the inventory taxonomy should align with recognised definitions and standards to improve clarity and comparability across systems.

10.18 Banking institutions must ensure that only AI systems recorded in the inventory are allowed to be deployed into production. All change management records and approvals must be linked to the relevant inventory entry.

PROPORTIONALITY AND RISK-BASED APPLICATION:

The principles in Sections 11 to 16 are not intended to be applied uniformly across all AI systems. Banking institutions must implement these principles in a manner proportionate to the nature, size, complexity and risk profile of the institution, as well as the risk tier, materiality, intended use and level of autonomy of the relevant AI system.

11. DATA GOVERNANCE AND QUALITY

11.1 Banking institutions must maintain an institution-wide data governance framework for AI that defines how data is collected, prepared, used, and controlled across the AI lifecycle.

Public

11.2 Clear data roles and responsibilities must be assigned, including accountable data owners and stewards for each dataset used to train, validate, test, or operate AI systems.

11.3 Banking institutions should maintain appropriate records of material datasets used in the development, validation, testing and operation of AI systems. The level of documentation should be proportionate to the risk tier and materiality of the AI system.

11.4 For third-party AI systems, institutions should maintain information that is reasonably available to them from service providers, including information obtainable through contractual arrangements, vendor documentation, independent assurance reports and other reasonable due-diligence processes.

11.5 Banking institutions must define and apply measurable data quality standards. These must include measures such as accuracy, completeness, timeliness, and consistency, with acceptance thresholds documented for each dataset.

11.6 Data preparation activities should be standardised and fully documented. This includes labelling, annotation, cleaning, enrichment, and aggregation, to support auditability and reproducibility.

11.7 Before model training, the suitability and representativeness of each dataset must be assessed for its intended use. Any assumptions, coverage limitations, or known gaps must be clearly documented, taking into account the data provenance, including the origin and method of collection of the dataset.

11.8 Banking institutions must test datasets for unwanted bias across relevant groups. The results must be documented, and appropriate bias mitigation measures must be implemented before model training.

11.9 Banking institutions must implement documented bias mitigation measures where bias is identified and ensure that such measures are monitored and reassessed on an ongoing basis throughout the AI lifecycle, including during model deployment and operation, commensurate with the risk tier of the AI system.

11.10 Training, validation and testing datasets must be strictly separated and used only for their intended purpose to ensure reliable model performance.

11.11 Banking institutions must maintain data lineage and version control for all AI datasets. Original datasets must be archived in accordance with applicable legal, regulatory, contractual, and data-retention requirements, and records must clearly link each model version to the specific dataset versions used.

11.12 Data retention and deletion rules must align with business needs and applicable laws and regulations. Banking institutions must define and document retention periods for all AI datasets, considering the intended use, risk tier, legal obligations, and audit

Public

requirements. Evidence of secure data disposal must be maintained when datasets are no longer required.

11.13 Access to datasets must follow the least privilege principle, meaning that access is limited to what is necessary. All approvals for internal or external data sharing must be properly documented.

11.14 Datasets obtained from third parties should undergo documented due diligence. This includes an assessment of legal rights, data quality, data provenance, and suitability for the intended use.

11.15 Banking institutions must conduct periodic reviews of datasets to identify data drift, quality deterioration, or reduced relevance. Decisions to retain models or refresh data should be based on the risk tier of the AI system.

11.16 Banking institutions must be able to provide comprehensive data documentation to internal audit and the Bank upon request. This should include the dataset register, data origin records, quality assessments, bias testing results, and data lineage logs.

12. TRAINING, TESTING AND MONITORING OF AI SYSTEMS

12.1 Banking institutions must ensure that AI systems are trained using data that is relevant, representative and appropriate for the intended purpose.

12.2 Training processes must document key assumptions, identified limitations, and measures implemented to address bias and data quality risks, commensurate with the system's risk tier.

12.3 Before deployment to production, banking institutions must conduct and document a pre-deployment assessment covering the intended use case, material risks, data considerations, human oversight arrangements, validation outcomes and approval status. Documented acceptance criteria, test plans and independent validation proportionate to risk must be completed and approved.

12.4 Testing must assess model performance and robustness under both normal and adverse scenarios, including error handling and stability across data shifts.

12.5 Fairness tests and explainability assessments must be conducted at a level appropriate to the system's risk tier, with results documented and any mitigation measures implemented and verified.

12.6 An AI system must only proceed to production once there is clear evidence that acceptance criteria have been met and all required approvals have been obtained in accordance with the institution's governance process.

Public

12.7 During production, ongoing monitoring must track performance, model drift, bias and fairness, and exceptions with thresholds that trigger human review, temporary suspension, rollback, or retraining.

12.8 Monitoring results, including bias and fairness assessments, interventions, and corrective actions must be documented, periodically reviewed, and used to support continuous improvement.

12.9 Any material changes to the model, data, integration, or use case must trigger re-validation and re-assessment of risk prior to deployment.

12.10 Material incidents must be promptly escalated, investigated, and, where required, reported to the Bank within 24 hours of the banking institution becoming aware of the incident. Near-misses must be recorded and analysed to support learning and risk mitigation.

13. EXPLAINABILITY AND INTERPRETABILITY OF AI SYSTEMS

13.1 Banking institutions must ensure that AI systems are designed and operated with risk-based, proportionate explainability aligned to the risk tier and materiality of the AI system.

13.2 Explainability must enable relevant stakeholders, including users, internal control functions, and the Bank, to understand how the AI system operates, the key factors influencing outputs and the limitations of the model.

13.3 For high-risk AI systems, explanations must be sufficiently robust to support human review, challenge, auditability, and justification of outcomes, particularly where decisions materially affect customers, employees or regulatory compliance.

13.4 Banking institutions must implement appropriate methods and controls to support explainability, including model documentation, traceability of inputs and outputs, and identification of key decision drivers.

13.5 Explainability must be assessed during the validation and approval process prior to deployment and monitored on an ongoing basis thereafter. Any material changes to the AI system must trigger a reassessment of explainability.

13.6 Banking institutions must ensure that high-risk AI systems are designed and selected to enable interpretability, including the ability to understand the underlying model structure, input-output relationships, and key assumptions.

13.7 Banking institutions must ensure that the use of complex or less interpretable AI models is appropriately justified and subject to enhanced controls, including stronger

Public

validation, monitoring, and human oversight, commensurate with the risk tier and materiality of the AI system.

14. SECURITY, DATA PRIVACY, CUSTOMER CONSENT AND PROTECTION

14.1 In line with the requirements of BID-30 and applicable national data protection and privacy laws, banking institutions must implement AI-specific cybersecurity controls to address threats such as data or model poisoning, prompt injection, model manipulation or theft, and content-related harms.

14.2 Where personal data is processed, banking institutions must ensure that a lawful basis for processing is clearly established, consent is obtained where required, and appropriate notices are provided, including disclosures relating to automated decision-making where necessary.

14.3 When customers interact with AI systems, banking institutions must clearly disclose the use of AI, provide meaningful information on the key factors influencing material customer outcomes, inform affected customers of the availability of human review or intervention, and ensure that such disclosures are clear and proportionate to the nature, materiality and potential impact of the AI use.

14.4 Banking institutions must apply privacy-preserving techniques such as pseudonymisation or encryption, where appropriate, to support data use for testing, evaluation and fairness monitoring while protecting individuals.

14.5 Banking institutions must establish effective complaint and appeal processes that allow customers to raise concerns, submit additional information, and obtain timely human review of material decisions.

14.6 Banking institutions must establish and implement appropriate controls governing the use of GenAI tools by employees and third parties. Such controls must include acceptable use requirements, restrictions on the input of sensitive or confidential information, monitoring of usage where appropriate, and integration into existing information security, data governance or acceptable use policies.

14.7 Where AI systems are provided by third parties and direct access to underlying model information is limited, banking institutions shall take reasonable and proportionate steps to obtain sufficient assurance regarding the system's performance, security, governance and risk profile, having regard to the materiality and risk of the AI system. Such assurance may be obtained through contractual arrangements, independent assurance reports, certifications, audits, or other appropriate due diligence measures.

15. ETHICAL USE OF AI SYSTEMS

Public

15.1 Banking institutions must ensure that AI systems are designed, developed, and used in a manner that is trustworthy and aligned with recognised ethical principles, including fairness, transparency, accountability, and respect for human rights and human dignity.

15.2 AI systems must be used in a manner that does not result in harm, including unlawful discrimination, unfair bias, or disproportionate adverse outcomes, particularly where such systems affect customers, employees, or other stakeholders.

15.3 Banking institutions must not deploy or use AI systems that are deceptive, manipulative, or exploitative, including systems that mislead users, exploit behavioural vulnerabilities or unduly influence decision-making in a manner that could cause harm.

15.4 AI systems must be subject to appropriate human oversight, such that material decisions can be reviewed, challenged, and, where necessary, overridden.

15.5 Banking institutions must establish controls and processes to identify, assess, and mitigate ethical risks arising from the use of AI systems throughout the lifecycle, in line with their risk management framework.

15.6 The Board of Directors and Senior Management remain ultimately accountable for ensuring that AI systems are used in a responsible, ethical, and lawful manner, including where such systems are developed or provided by third parties.

16. SUPERVISORY NOTIFICATION AND INFORMATION REQUIREMENTS

16.1 Banking institutions must comply with the following requirements:

a) Maintain a single, consolidated AI inventory in accordance with paragraphs 10.1 to 10.4 of these Guidelines. The inventory must be kept up to date, be subject to periodic review as part of the institution's assurance and oversight processes, and be made available to the Bank upon request, including during supervisory reviews and onsite examinations.

b) Notify the Bank within 24 hours of becoming aware of a material incident, as defined in these Guidelines. Where complete information is not immediately available, institutions may submit an initial notification containing preliminary information and shall provide supplementary updates as additional information becomes available. The Bank may request additional information or supporting documentation to facilitate its assessment of the incident.

Public